



A11106 430902

NIST
PUBLICATIONS

NISTIR 6536

High Speed Network Applications and Implications for Fiberoptic and Copper Connections

Dean Collins
John Antonishek
Sean Sell
Alan Mink

U. S. DEPARTMENT OF COMMERCE
Information Technology Laboratory
National Institute of Standards
and Technology
Gaithersburg, MD 20899

June 2000

**NIST**

**National Institute of Standards
and Technology**
Technology Administration
U.S. Department of Commerce

QC
100
.U56
NO.6536
2000 c.2

High Speed Network Applications and Implications for Fiberoptic and Copper Connections

**Dean Collins
John Antonishek
Sean Sell
Alan Mink**

U. S. DEPARTMENT OF COMMERCE
Information Technology Laboratory
National Institute of Standards
and Technology
Gaithersburg, MD 20899

June 2000



U.S. DEPARTMENT OF COMMERCE
Norman Y. Mineta, Secretary
TECHNOLOGY ADMINISTRATION
Dr. Cheryl L. Shavers, Under Secretary
of Commerce for Technology
NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
Raymond G. Kammer, Director

High Speed Network Applications and Implications for Fiberoptic and Copper Connections

Dean Collins, John Antonishek, Sean Sell, Alan Mink
National Institute of Standards and Technology
100 Bureau Drive, Stop 8950
Gaithersburg, MD 20899-8950

Abstract

We present an analysis of the current tradeoffs between using copper wire versus fiberoptic cable for rewiring buildings at NIST for data communications to the desk-top. We consider tradeoffs from two different viewpoints: (a) the cost of installing and maintaining the wiring, and (b) the computer imposed limitations on the use of the wiring. These findings substantiate our decision to use copper wiring.

Introduction

The National Institute of Standards and Technology (NIST) is presently upgrading the wiring of existing buildings and also installing wiring in a new building presently under construction. This new building is a modern laboratory. The goal of the building wiring upgrade is to provide NIST with as modern a data communication network as possible within cost constraints. We also have an upgrade to an Asynchronous Transfer Mode (ATM) fiberoptic backbone under way, which will not be addressed here.

NIST is a part of the Department of Commerce. At the Gaithersburg, Maryland campus it covers 2.34 km^2 (≈ 578 acres) and has 31 buildings serving a staff of over 3000. The campus buildings are connected with a fiberoptic Fiber Distributed Data Interface (FDDI) backbone. The local wiring within the buildings is mostly Category 3 (CAT 3) copper wiring (16 Mbit/s) with selected sites using Category 5 (CAT 5) (155 Mbit/s) copper wiring and fiberoptic connections. Most of the building wiring runs are less than 100 meters from the network rooms.

The rewiring of existing buildings and wiring of the new laboratory involve tradeoffs between copper and fiberoptics and the associated switching devices. Fiber optic cable comes in two general types, multi-mode and single-mode. Single-mode fiber is a bit more expensive (about 5 %) than multi-mode, but associated Network Interface Cards (NICs) and switches can cost significantly more. Single-mode fiber is capable of supporting higher bit rates over longer distances than multi-mode, but is not necessary for the desk-top environment. For the remainder of the paper our use of fiber optic cable should be taken to mean multi-mode fiber. We looked at the various tradeoffs from two different viewpoints: (a) cost of installing and maintaining the wiring, and (b) computer imposed limitations on the use of the wiring. The following discussion presents our findings as we compare copper and fiberoptic connections to the desk-top for our use at NIST. These findings substantiate our decision to use copper.

Wiring Considerations

Our major wiring options included CAT5 Copper, Enhanced CAT5, Gigabit Copper, and Fiberoptics. The data rates and run lengths possible at these data rates are listed in Table 1. Copper wiring is limited to a 100 m solution. In a typical NIST office building, this constraint does not present any undue hardship in the location of the network rooms. With fiber, it is possible to service more, if not all, of the building from one central network room location, but in practice this becomes cumbersome because of the amount of space necessary to support all of the cable and the limited amount of space available on a cable tray.

	Data Rates (Mbit/s)	Run Length (m)
CAT5	155	100
Enhanced CAT5	622	100
Gigabit (Gigaspeed)	1000	100
Fiberoptic	>1000	>300

Table 1. Wiring Performance

Network active devices are about four to six times more expensive, per port, for fiber optic cabling than for copper wiring. The density of fiber optic ports is also much lower and requires additional chassis and modules. Since the fiber optic equipment supports a lower port density, more patch panel and equipment space is necessary (almost 4 times the equipment space and one and a half times more patch panel space). Since fiber requires more chassis to support the same amount of ports, it also requires additional air conditioning. Furthermore, fiber patch cords used in the network closets and offices are five times more expensive than their copper counterparts.

The installation costs for fiber (including pulling, terminating, testing, etc.) are twice that of copper wiring. Since maintenance agreements are a percentage of the initial cost, the same factors for component costs and installation costs are relevant. More in-house staff are trained to pull, terminate and test copper wiring versus fiber optic cabling, thus reducing the need for external support for maintenance of a copper wiring infrastructure.

Let us now compare the differential cost of wiring up a typical NIST laboratory considering only the installation and component costs. The cost difference, shown in Table 2, for one laboratory, is approximately \$1,200,000. If this is spread across the entire campus, it is our estimate that the differential additional cost for using fiber optics would be approximately \$14,500,000.

	Copper (CAT5)	Fiber
Installation	\$327,000	\$700,000
Components	<u>\$234,759</u>	<u>\$1,064,655</u>
Total	\$561,759	\$1,764,655

Table 2. Installation and Component Costs for Fiber vs. CAT5 Copper

Based on our evaluation we have reached the following conclusions. For rewiring the existing buildings we chose CAT5 copper, which provided cost benefits and sufficient capacity, of up to 155 Mbit/s, to support the majority of network users. This particular decision was based on knowing that in five to ten years all the existing buildings are planned for renovation and thus a chance to reevaluate wiring requirements and options. For the wiring of the new laboratory we chose Gigabit copper because it is the most advanced copper solution currently available and there is a minimum cost differential for the material but no additional installation cost. It is only 10 % more expensive than CAT5, and 5 % more expensive than using enhanced CAT5. Furthermore it is completely backwards compatible with existing CAT3 and CAT5 installations.

Computer Considerations

Let us now take a look at the issue from a different standpoint. The computer user does not care about electrons or photons. What the user is interested in is performance (meaning effective communication latency and bandwidth), interoperability, availability, ease of use, and the cost inherent in the computer itself. Simply put, what measurable benefit will the computer user obtain for the extra cost of fiberoptics? In order to answer this question, let us again look at our four options: Category 5 copper, Enhanced CAT5, Gigabit copper, and Fiberoptics. From the standpoint of interoperability and ease of use there is no real advantage for either copper or fiber. The transmission media (copper vs fiber) is transparent to the user, even to the system device driver programmer. The NICs hide that level of detail. The device driver software in the operating system deals with a logically delimited data packet (or frame, or segment, etc.). Therefore, there is no difference in ease of use, or for that matter use!

Interoperability of the medium is handled by the NIC and network device, and is transparent to the user. Currently all processing is done electronically in parallel (8, 16, or 32 bits) for two reasons: (1) the data are meaningful only in groups and by dealing with a parallel group, one has more time to process the data than at the bit rate; (2) optical processing of protocols is not yet feasible, but research is progressing in this area. [1, 2] Interoperability of protocols is handled by the protocol stack. This allows only well known protocols to be used together. For example, routers convert between a number of well-known protocols such as Ethernet and FDDI. Non-standard protocols (e.g., many Applications Programming Interfaces (APIs)) cannot interoperate! Well

known standard protocols also provide stability and reliability, mostly through massive exposure and fine tuning over an extended time period. Performance on the other hand is achieved via nonstandard APIs or O.S. (Operating System) Bypass protocols, although an API called VIA (Virtual Interface Architecture) is being proposed by an industrial group, which includes Intel. [3, 4]

Now let us consider the limitation to communication speeds. Figure 1 compares memory access speeds (CPU-to-memory) with network speeds (Network I/O-to-memory). The Y axis shows common network speeds. The X axis shows memory access speeds. Two representative points on the X axis are PCs based on a 200 and 400 MHz Pentium microprocessor. The memory access (reads and writes) speed is




Figure 1. Limitations to Communication Speeds

derived by doubling the measured memory copy (read-followed-by write) speed. Network communication requires user data to be encapsulated and transferred from memory (usually specified by its virtual address) to a NIC plugged into the I/O bus. The actual transfer is accomplished via direct memory access (DMA) by the NIC. The NIC “packages” the data and transmits it onto the communications media, possibly traversing through a number of switches (which may buffer and process some or all of the message before passing it along). When it reaches the destination NIC, the

process is reversed, placing the data into memory (a user-specified buffer).

The dashed diagonal line in Figure 1 represents the upper limit of communication speed, equal to memory access speed. This upper limit may be a desirable goal, but not practically achievable.

All current microprocessors running TCP/IP can use the maximum bandwidth of the current, common commodity networks (Fast Ethernet & ATM/OC3). Current processors can execute the instructions for the common TCP/IP protocols faster (somewhere between 150 Mbit/s and 600 Mbit/s) than the bandwidths of these networks. Our performance measurements [5] have shown that although ATM/OC3 has a 50 % higher bandwidth than Fast Ethernet, using TCP/IP over ATM/OC3 only obtains about 30 % higher communications throughput than TCP/IP over Fast Ethernet. Furthermore, on a 200 MHz PC, TCP/IP would limit communications throughput to less than 200 Mbit/s. Thus, providing a fiberoptic Gbit/s network interface for such a machine would result in only marginal communications throughput compared to an OC3 link.

However, to achieve Gbit/s communications (Gigabit Ethernet and ATM/OC12) requires "Operating System (O.S.) Bypass" protocols. These significantly reduce the intervention of the processor to execute the communication protocol by moving most of the previous software instructions to hardware. This technique is limited by the speed of the I/O bus. In many computers today that is the PCI bus. To achieve O.S. Bypass protocols requires that the communication be done directly to the user buffers, bypassing the O.S. This requires that the user buffers be locked in physical memory and that the NIC be programmable to maintain the list of user buffer physical addresses (page tables) and be able to multiplex different message streams from different sources. To achieve even higher speed communication requires faster networks and switches connected to faster I/O busses, or even direct attachment to the memory bus, in addition to bypassing the O.S.

When considering communication latency for local area networks, it turns out that network latency is not a determining factor in selecting either copper or fiber. Latency depends on the effective communication speed and the length of the packet being sent. Consider the latency of a 200 MHz Pentium running on a Fast Ethernet LAN shown in Figure 2. For this comparison we have chosen two cases: (a) minimum data packet length of 4 bytes and (b) maximum data packet length of 1453 bytes. Depending on the packet size used, the communication latency can range from $\approx 84 \mu\text{s}$ to $\approx 248 \mu\text{s}$. The protocol stack terms are dependent on processor speed and could also be reduced by advanced protocols such as O.S. Bypass. The NIC DMA transfer terms are only effected by memory access speeds. It should be noted that the only term dependent on bandwidth is the Network Latency term, a function of network speed, length and packet size. If a gigabit connection was used, the latency associated with the maximum packet length would decrease from approximately 122 to 12 microseconds. However, the total communication latency would still be dominated by the non-bandwidth related terms.




Figure 2. Communication Latency Breakdown for a 200 MHz Pentium on a Fast Ethernet LAN

Looking at cost and availability of NICs and their associated device drivers shown in Table 3, we see that Fast Ethernet is a defacto standard and thus inexpensive. A device driver is a piece of software, written by a systems programmer, that is part of the operating system and acts as an interface between a specific device and the operating system. On the device side, this piece of software handles all the specific details of a single device, such as an NIC, a disk, a keyboard, etc. This software is so specific that even different models and/or versions of a device, from the same manufacturer, may require a different device driver. On the operating system side, this piece of software hides all those device specific details and presents a uniform interface to the operating system for that class of device. An ATM/OC3 NIC is available for most major machine I/O busses and is now moderately priced, but device drivers are not available for all operating systems. The faster ATM/OC12 NIC is very sparsely available and so are the necessary device drivers; it is also expensive. The new Gigabit Ethernet is just starting to become available. These early offerings are expensive, but will undoubtedly become cheaper as it matures. Very fast ATM connection (above OC12) are not currently aimed at the desktop and NICs are not

available.

	<u>Network Interface Cards</u>	<u>Device Drivers Software</u>
Fast Ethernet	Inexpensive ≈\$60	De facto Standard
ATM/OC3	Moderate ≈\$550	Selectively Available
ATM/OC12	Very Expensive ≈\$1000+	Very Sparse to N/A
Gigabit Ethernet	Expensive ≈\$1000	Selectively Available (early)
ATM/OC24 & 48	N/A	N/A

Table 3. Cost and Availability of NICs and Associated Device Drivers

From the standpoint of the computer user, the choices are dependent on their selection factors.

- If price/performance & interoperability are the most important factors, Fast Ethernet is the choice. It has low cost, uses copper, has good performance, high availability, and is the de facto standard, thus providing interoperability.
- If local performance is the only important factor, Gigabit Ethernet or ATM-OC12 or Myrinet are good choices. They are fast networks, use fiber (or copper for Gigabit Ethernet), moderate to high cost, selective availability, but sparse standards.
- If global interoperability and performance are the most important factors, use either Fast Ethernet or ATM-OC3. These have good performance, use copper and/or fiber, have moderate to low costs, high to reasonable availability, support standard IP protocols for interoperability, and have long distance capability.

Conclusions

At the present time, for our NIST application, copper is still the clear winner. On wiring considerations alone, copper provides NIST definite cost benefits while supporting sufficient bandwidth in our five to ten year horizon. NIST computer users cannot utilize the longer run length and higher bandwidth of fiber, thus NIST users would see no performance difference between copper and fiber. Fiberoptic technology is changing rapidly, however, so is copper technology. In order for fiberoptic technology to become attractive for the particular scenario we have looked at, some subset of the

following must occur:

- Components and NICs must become more readily available and cost less.
- Device drivers must be readily available and supported.
- Fiber must have a smaller footprint to allow high densities on patch panels and network devices.
- Fiber must become easier and less expensive to terminate.
- Bandwidth needs must require fiber.
- Standard protocols other than TCP/IP must be developed and utilized.

References

- (1) R. Bencek, et al., "1.24416 Gbit/s Demonstration of a Transparent Optical ATM Pocket Switch Node," Electronics Letters, Vol. 30, No. 7, 31 Mar 1994, pp 579-580.
- (2) I. Glesk, J. Solokoff and P. Prucual, "All-Optical Address Recognition and Self-Routing in a 250 Gbit/s Pocket-Switched Network," Electronic Letters, Vol. 30, No. 16, 4 Aug 1994, pp 1322-1323.
- (3) D. Dunning, et al., "The Virtual Interface Architecture," IEEE Micro, Mar-Apr 1998, pp 66-76.
- (4) T. von Eicken and W. Vogels, "Evolution of the Virtual Interface Architecture," IEEE Computer, Nov 1998, pp 61-68.
- (5) M. Indovina, A. Mink, R. Snelick and W. Salamon, "Performance Measurement of ATM and Ethernet Computing Clusters," Proc ATM98 Development Int'l Conf/Exhibits, Mar. 1998, Rennes, France, Mar. 1998, pp 45-64.

Dean R. Collins, Chief, High Performance System & Services Division
John K. Antonishek, Group Leader, Network & Telecommunications Systems Group,
High Performance Systems & Services Division
Sean Sell, Project Leader, Network & Telecommunications Systems Group, High
Performance Systems & Services Division
Alan Mink, Project Leader, Scalable Parallel Systems & Applications Group, High
Performance Systems & Services Division
Contact: Mary Floyd, mmfloyd@nist.gov, Fax 301-963-9137, Phone 301-975-2869

